Remarks

1. INTRODUCTION.

The Examiner rejected claims 1-6 under Section 103(a) as being unpatentable over Crockett in combination with Tzelnic et al. ("Tzelnic"). The Examiner posits essentially that all the limitations of claim 1, labeled (a) through (m) for identification, are taught by Crockett except for limitation (l), which the Examiner contends is taught by Tzelnic. Applicant respectfully traverses the rejections and requests reconsideration for the reasons discussed below. The Examiner has not shown a *prima facie* case of obviousness.

In this response to the non-final rejection, claim 1 amended merely to clarify that the method calls for signaling an operator, rather than an application program, when closing the source data set is requested. New claims 7-55 are added to more completely claim the invention. In the specification, paragraph 11 is amended to correct minor errors.

2. SUBSTANCE OF THE INTERVIEW IN THE PTO.

The May 18, 2006 interview was attended and made of record by Examiner Ko and Primary Examiner Bragdon. The inventor, Mr. Perego, and the undersigned attorney for applicant attended. We discussed how Crockett et al. teaches volume level migration as opposed to data set level migration as set forth in applicant's claims. For example, it was noted that a data set level migration requires first selection and then an initial assessment of the selected data set(s) specified to be migrated ("source data set"). This initial assessment entails determining the locations of all the extents of the source data set –on what volumes and tracks they are located –so that they can be copied. (See note 4.) This assessment also is used to make an appropriate initial allocation of the target data set. In a volume level migration, by contrast, as described in Crockett et al. and known commercial products, the system need not and does not locate specific data set extents; rather, it simply copies entire volumes, without regard to what particular data sets may or may not be involved.

Further, at the interview, there was discussion of aspects of the invention related to changes made to the source data set allocation after the initial assessment was made and copying started. The source data set can change during the mirroring process, e.g., it might be extended, or space released, or the data set removed entirely. In this regard, step K of original claim 1 was noted as calling for accommodating, during the downtime

window, any allocation changes affecting the source data set since the initial assessment, so that necessary changes can be made in the target data set allocation. Because Crockett et al. do allocate target data sets at all (again, they operate at the volume level), these features are not disclosed in the prior art. The Examiner requested discussion of these difference in the present response; they are discussed below.

3. THE CLAIMED INVENTION AS A WHOLE – *DATA SET LEVEL* MIGRATION, WOULD NOT HAVE BEEN OBVIOUS FROM THE PRIOR ART *VOLUME LEVEL* COPYING.

To establish a *prima facie* case of obviousness, as the Examiner is aware, three basic criteria must be met. First, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings. Second, there must be a reasonable expectation of success. For the reasons discussed later, the Crockett and Tzelnic references actually are incompatible, and thus one cannot expect success in combining them.

Moreover, the prior art references when combined, if they could be combined, must teach or suggest all the claim limitations. The teaching or suggestion to make the claimed combination and the reasonable expectation of success must both be found in the prior art and not based on applicant's disclosure. *In re Vaeck*, 947 F.2d 488, 20 USPQ2d 1438 (Fed. Cir. 1991). MPEP 706.02(j). For the reasons below, even if the references were properly combinable, *arguendo*, the combination would not render the claimed invention obvious. Indeed, almost none of the claim limitations is disclosed in the prior art cited, because of the fundamental difference between data set-level and volume-level storage system operations.

A. Discussion

1. As discussed in the interview, prior art mirroring solutions, like those disclosed in Crockett et al. ("Crockett"), are volume-oriented processes. That is, they copy or backup DASD storage volumes, track by track, without regard to the actual data (and or meta-data) stored on those volumes. This is known as volume-level or volume-centric mirroring. Examples include the well-known PPRC and XRC VMS utilities. Claim 1, as discussed below, is directed to migration of a specific data set (called a "source data set"),

rather than a volume.[1] Accordingly, in the method of claim 1, tracks that are not part of the source data set are not copied; they are skipped.

2. Because the method of claim 1 ignores (does not copy) tracks that do not store a portion of the selected data set, the method of claim 1 does not require monitoring all changes to the primary volumes. Rather, only changes to the tracks identified as storing a portion of the data set being migrated must be monitored for changes to maintain synchronization.

3. Crockett, PPRC, and the like do not employ a "track translate table" as required by claim 1. Indeed, they would have no use for such a table, because again, in volume-centric mirroring, every track is simply copied to the corresponding track location on the secondary. Thus the CCH locations are the same as on the source or primary volume. Accordingly, the primary and secondary DASD units in the prior art mirroring must be identical. An advantage of the method of claim 1 is that it can accommodate dissimilar storage devices, implying different CCH locations. Thus it is especially useful for migration of data sets to larger volume storage systems. This feature, however, does require changing the meta-data (catalog data), as discussed later.

B. Specific Limitations of Claim 1.

For at least the foregoing reasons, the teaching of Crockett taken as a whole, with or without the addition of Tzelnic, is dramatically different from the invention of claim 1. That said, a more specific discussion of some of the individual limitations will serve to highlight the differences.

Referring now to claim 1, the Examiner pointed to various passages in the prior art alleged to disclose the claim limitations. Applicant respectfully responds as follows.[2] For each limitation, the subject claim language is reproduced first in Arial font, followed by the passage(s) in Crockett (and later Tzelnic as to one limitation) that were cited by the Examiner, in courier font, and then finally Applicant's remarks in Times New Roman.

---

[1] Hence the present title: "*Data Set Level Mirroring...*"

[2] In the interest of brevity, Applicant does not discuss every claim limitation. Because the prior art must teach or suggest *all* the claim limitations to negative patentability, the absence of even one limitation would suffice to avoid anticipation.

"1. (original) A method for combining extents of a source data set, with minimal down time imposed on applications accessing the source data set, the source data set having a corresponding name, source DASD device(s) and one or more extent locations on the source DASD device(s), the method comprising the steps of:

a) assessing the source data set to be migrated, the volume(s) it resides on, and the total space it occupies;"

The Examiner cited Crockett column 7, lines 65-67; column 8 lines 13-16. Those passages read:

"This consists of reading a predetermined number of tracks from the primary and then writing them out to a secondary, and repeating this read and copy operation over the entire volume...

"One implementation could consist of computing the addresses of a predetermined number of consecutive primary..."

As mentioned earlier, Crockett discloses reading and then writing tracks, in other words, copying tracks, "and repeating this read and copy operation over the entire volume." (Emphasis added, column 7, line 67 –column 8, line 1.) Applicant's Claim 1 above refers to a specific "source data set" and calls for "assessing the source data set to be migrated, the volumes it resides on," etc. Crockett just copies whole volumes, without assessing data sets at all.[3] Thus this limitation is not disclosed.

"b) allocating space for a corresponding target data set using the total source space as a primary allocation request and using the original source data set name;"

The Examiner cited to Crockett column 8, lines 16-24, and did not elaborate. That passage reads as follows:

[3] The Background section recites: [0006] "Mainframe files or "data sets," for example under MVS operating systems, are allocated space on DASD devices in one or more contiguous groups of tracks. (A typical DASD device has 15 tracks per cylinder.) Each contiguous group is called an "extent." The number of permissible extents for a data set varies by type of data set, level of the OS, and other factors, but generally some fixed limitation is imposed. A data set can span more than one volume, although an individual extent must be stored in a single volume. This presents two common problems: a) a data set cannot expand because the maximum number of extents has been reached (even though DASD space is available), and b) due to space available at the time an extent was needed, a data set may be spread over more volumes than the user would like."

"...reading the tracks and establishing them in step 309, and sending them to the secondary SCU in step 311. This repeated via a branch 313 back to step 305. There, the next group of track addresses is determined again in step 305. In this case it would be addresses for tracks 3, 4, and 5. This process would be repeated until all of the tracks on the primary volume had been copied and sent to the secondary SCU. The term "establish", as used in step 307 means that after each group of three tracks is read, ..."

Again, the reference itself describes a volume-centric methodology: "This process would be repeated until all of the tracks on the primary volume had been copied and sent to the secondary SCU." The method of claim 1 is more selective. It first analyses the specific data set to be migrated (step (a)), and then allocates space for a corresponding target data set .... Crockett says nothing about allocating space based on data set requirements. Rather, Crockett deals with entire volumes, i.e. no need nor concern with data sets on volumes being copied. The referenced Crockett text describes how groups of tracks are involved in the initial copy process. These groups of tracks have no relationship to data sets on the volume(s) being copied. Thus this limitation is not disclosed.

"c) designating one or more target locations in the allocated space and assigning each source data set extent location to a respective one or more of the designated target locations;"

The Examiner cited to Crockett column 10, lines 8-12. That passage reads as follows:

"Prior to writing out the record sets, SCU 31 ascertains whether the address of the tracks read during volume initialization and so far written to the secondary volume 33 has proceeded past the highest address for the record sets."

Applicant's claim limitation (c) speaks to initially finding and then allocating locations on target volumes for data sets on source volumes to be migrated. Specifically, it calls for designating target locations and "assigning each source data set extent location ....". Crockett does not assess data sets, as noted above, and thus is not aware of any data set extent locations.[4] The Crockett passage quoted above deals with whether records sets (groups of continuous source record updates) fall below or above the records copies so far

[4] An "extent" is a contiguous group of tracks (of a DASD) allocated to a data set. A single data set often comprises more than one extent, and the extents may span multiple volumes. See applicant's FIG. 2.

during what Crockett calls the monotonic copy (initial sequential copy of all tracks on a volume). Importantly, "record sets" are not "data sets," as Crockett explains in his patent:

> "More particularly, the first step of the method, namely that of initially synchronizing the tracks of data on the primary DASD volume with counterpart tracks on the secondary DASD volume, comprises several substeps. These substeps include reading from the primary volume of a predetermined number of tracks as a group in a monotonic address order and copying said tracks in that address order on the secondary volume. The substeps further include forming record sets of CPU-originated updates to the tracks on the primary volume and copying those record sets to the secondary volume having addresses less than the highest address of the primary tracks copied onto the secondary volume." [Column 5, lines 14-25. Emphasis added.]

For these reasons, it is apparent that limitation (c), including assigning each source data set extent location to one or more target locations, is not disclosed or suggested by Crockett.

"d) starting monitor programs on any images that can write to any of the identified source volumes in order to detect a subsequent change to the source data set;"

Regarding this claim limitation, the Examiner cited to Crockett column 8, lines 23-26. That passage actually reads as follows:

"The term "establish", as used in step 307, means that after each group of three tracks is read, the primary SCU is enabled to monitor the occurrence of any application updates to that group of primary tracks."

Note that Crockett enables the primary SCU (Storage Control Unit) to monitor for updates (changes) to a group of tracks after is it read. The SCU is of course a hardware device (e.g. the IBM 3990 shown in Crockett FIG. 2),[5] which can be employed to

[5] Crockett explains in the description of related art: "When an application runs on a multiprocessing CPU, such as an IBM S/390 with an MVS operating system, it will generate read or write calls for data to the operating system (OS). If the data is not present in CPU main memory, the OS will invoke an access method and establish a path to the data. The path will lead to data stored or to be written on one or more DASDs in an attached storage subsystem. The storage subsystem may be of the demand/responsive,

implement the monitoring described by Crockett. Claim 1(d), to the contrary, calls for "*starting monitoring programs* on any images that can write to any of the identified source volumes..." thus the claim calls for monitoring via software, not the SCU hardware. The Specification says, "In general, aspects of the present invention will be embodied in a software program or utility we will call an implementation or a solution." Paragraph [0035].

"e) storing an indication of each source data set track detected by a monitor program as having changed;"

Notice this limitation does NOT call for storing indicia of any track changed. Rather, it calls for storing indicia of only the source data set track changes; i.e., changes to only those tracks that are part of the particular data set of interest. Once again, the claim is focused at the data set level. Moreover, as noted above, this describes software solution (employing a "monitoring program"). Crockett is dependent on a combination of hardware/microcode and system software such as the SDM (system data mover) whereas the monitor program of claim 1e has the advantage of removing any DASD manufacturer requirements and any reliance on the SDM.

"f) begin copying the source data set in accordance with said assignment of each source data set extent so as to form the corresponding target data set;"

The claim calls for copying the source data set. Again, one of many distinctions over the prior art is that Crockett copies all tracks of a volume, whereas this claim calls for copying only those tracks belonging to selected data sets. Second, the copying is "in accordance with said assignment of each source data set extent..." and thus copies to potentially *dissimilar* CCH (cylinder and track) locations are enabled. Under Crockett, source and target tracks must have identical CCH locations, and there is no "assignment of each source data set extent" to guide that copying as in claim 1.

---

hierarchically organized storage type. Illustratively, the IBM 3990 Model 6 storage control unit (SCU) is of that type. It includes a large multimegabyte cache, a nonvolatile store (NVS), and several redundant pathways to each of a plurality of 3390 DASDs or their equivalents." Column 1, lines 52-64.

C. Accommodating Source Data Set Changes.

Most, if not all, of the limitations of claim 1 can be distinguished over Crockett. However, recognizing that the Examiner has limited time, just a few additional points are discussed, including in particular the matter of accommodating changes to the source data set allocation subsequent to the initial assessment step of claim 1. The Examiner requested that this point be discussed in this response. In this regard, Claim 1 specifically calls for:

> "k) during the downtime window, accommodating any allocation ~~differences~~ changes affecting the source data set that occurred since the initial assessment of the source data set;"

A merge process in accordance with claim 1 initially copies only the tracks allocated to the selected source data set(s) at the beginning of the process. For the target to be a viable copy at the end of the process, which is the goal, changes in the status of source data sets during the mirroring process are accommodated. This is not disclosed in the prior art of record. Indeed, because of the several steps mentioned above, which are not shown in the reference, the system disclosed by Crockett can neither detect nor accommodate allocation changes to the source data set. The dependent claims specify some specific allocation changes: New source data set added (claim 4); source data set deleted (claim 5); extent changes (claim 6). Crockett does not copy selected data sets, only complete volumes; hence, it does not need to address this issue.[6]

D. Changing Catalog Entries.

Claim 1 further calls for:

> "l) changing catalog entries to reflect new target data set volumes..."

The Examiner correctly observed that, "Crockett fails to teach changing catalog entries to reflect new target data." The Examiner asserts, "Tzelnic teaches changing catalog entries to reflect new target data set (column 16, lines 1-5; column 18, lines 41-45)." Office action at page 4. Applicant respectfully dissents for these reasons: First, Tzelnic does not teach updating catalog entries "to reflect new target data set volumes" as

[6] The referenced portion of Crockett is dealing with write ordering of changed tracks to avoid back-leveling data, not anything to do with whether tracks copies are invalid because a source data was deleted, reduced in size, or extended.

is claimed here; and second, the two references are not properly combinable as suggested by the Examiner in any event because they are inconsistent, i.e., the teachings would be inoperable in combination.

The meta-considerations of combining Crockett and Tzelnic would not have been obvious because the methodologies are conflicting. Crockett is an improvement to mirroring technologies such as PPRC and XRC. The point in time image of a mirrored secondary is at the end of the process, or the split time. For the secondary to be usable, the meta-data representing the data on the mirrored volumes must be accurate for the PIT of the split. This is feasible, and in fact takes only nominal adjustment when entire volumes are mirrored, because the meta-data on a volume, i.e. the VTOC, VTOC Index, and VVDS, are included in total along with mirroring the data portion of the volume.

The problem remains, even applying the teachings of Crockett, that the correctness of the meta-data depends on manual user steps in the process if the PIT of the split is going to accurately represent the condition of the source at the same PIT. For instance, if a user does not close a data set on a volume being mirrored, meta-data like VSAM statistics and logical end of file indicators will not be entirely current.

Additionally, Crockett does not implement relocation of data, and thus it need not and does not address meta-data adjustments required because of relocation of data. As an example, extent differences that affect the VTOC contents and corresponding VVDS records are not an issue with Crockett because it is a volume centric solution. The present claim 1, for example, calls for "changing catalog entries to reflect new target data set volumes" to accommodate changes in data set locations.

Tzelnic teaches changing catalog and file system meta-data, not for merge-migration, but because the restore destination of a previously backed-up data set could be located somewhere other than the original location. However, the required changes with Tzelnic are simpler in nature because Tzelnic is working from a static copy of meta-data. Tzelnic admits that for a minimally disruptive backup as the source of subsequent restores, a technology like Snap is a preferred method for the backup. Snap, and equivalents, capture a PIT image of data at the beginning or time-zero of the process. Tzelnic alludes to capturing in some manner corresponding meta-data at the same time to facilitate restoring. Again, the present invention provides correct meta-data at the split time of the migration or mirror.

The meta-data captured at the beginning of a copy process is not valid, and cannot be used, for a PIT occurring at the end of a process. Moreover, the present invention in some embodiments permits previously existing data sets on mirror targets, which implies that the meta-data files cannot be copied in their entirety from the corresponding source volumes, and yet the pertinent meta-data records involved with the data sets being mirrored must be accurate, and accurate as of the end of the process, not a static copy obtained at backup time as in Tzelnic. For these reasons, the Tzelnic and Crockett methodologies are incompatible, and therefore one skilled in the art would not attempt to combine them. For these reasons, the Examiner has not made out a prima facie case of obviousness, and thus the rejection of Claims 1-6 should be reconsidered and withdrawn.

4. NEW CLAIMS 7-55 ARE ALSO PATENTABLE.

Only three of the new claims are independent. They read as follows:

"7. (new) A method for data set-level mirroring of a selected source data set while minimizing downtime of applications accessing the source data set, the method comprising:

identifying extents allocated to the source data set, each extent comprising a set of at least one DASD track;

allocating within a target data set at least a number of extents needed to store the identified extents of the source data set;

copying each track of an extent of the source data set to an extent of the target data set;

in response to any tracks of the source data set being changed by an application during the copying step, recopying the changed tracks from the source data set to the target data set;

commencing a downtime window in which applications may no longer access the source data set; and

in response to an allocation change being made to the source data set since the allocating step, performing a corresponding allocation change to the target data set during the down-time window."

The Examiner can readily observe that new claim 7 is clearly patentable over Crockett and Tzelnic at least because:

a. It is a method for data set-level mirroring of a selected source data set, not volume mirroring;

b. it calls for initial allocation based on the extents of the source data set;

c. it calls for copying each track of an extent of the source data set rather than simply copying whole volumes; and

d. it calls for, in response to an allocation change being made to the source data set since the allocating step, performing a corresponding allocation change to the target data set during the down-time window.

New claims 8-26 all depend from claim 7.

"27. (new) An article of manufacture comprising a computer-readable medium storing computer-readable program code for performing a method for data set-level mirroring of a source data set while minimizing downtime of applications accessing the source data set, the article of manufacture comprising:

computer-readable program code for identifying extents belonging to the source data set, each extent comprising a set of tracks;

computer-readable program code for allocating within a target data set at least a number of extents needed to store the identified extents of the source data set;

computer-readable program code for copying each track of an extent of the source data set to an extent of the target data set;

computer-readable program code for in response to any tracks of the source data set being changed by an application during the copying step, recopying the changed tracks from the source data set to the target data set;

computer-readable program code for commencing a downtime window in which applications may no longer access the source data set; and

computer-readable program code for performing an allocation change to the target data set in response to an allocation change being made to the source data set since the initial allocating step."

The Examiner can just as readily observe that new claim 27 is clearly patentable over Crockett and Tzelnic at least because:

a. It is directed to a computer program stored in a machine-readable medium that implements a method for data set-level mirroring of a selected source data set, not volume mirroring;

b. it includes computer-readable program code for identifying extents belonging to the source data set;

c. it includes computer-readable program code for copying each track of an extent of the source data set to an extent of the target data set rather than simply copying whole volumes; and

d. it includes computer-readable program code for performing an allocation change to the target data set in response to an allocation change being made to the source data set since the initial allocating step.

New claims 28-46 all depend from claim 27. Finally, new independent claim 47 reads as follows:

"47. (new) A method for data set-level migrating a selected source data set under an MVS operating system while minimizing downtime of the source data set, the method comprising:

(a) assessing the source data set to be migrated so as to identify a set of source data set tracks on which the data set is stored;

(b) storing initial meta-data that reflects extent locations of the source data set tracks;

(c) initially allocating a target data set with at least sufficient space to store the source data set tracks;

(d) beginning sequential copying of only the source data set tracks to the target data set;

(e) while the sequential copying proceeds, beginning monitoring to detect any changes to a source data set track after the track was last sequentially copied, thereby identifying as a modified track one that requires re-copying;

(f) continuing the sequential copying of source data set tracks until all of the source data set tracks have been copied, and all modified tracks have been re-copied;

(g) after completion of the sequential copying, requesting that the user close and unallocate the source data set;

(h) continuing said monitoring to detect and identify any modified source data set tracks until the source data set has been closed and unallocated;

(i) while waiting for the source data set to be closed and unallocated, re-copying any identified modified track;

(j) responsive to an indication that the source data set has been closed and unallocated, effecting a final re-synchronization step by re-copying any modified track not yet re-copied;

(k) checking to detect any allocation change to the source data set since initially allocating the target data set, by comparing current meta-data describing the source data set to the initial meta-data stored at the beginning of the process;

(l) in response to detecting an allocation change having been made to the source data set after the target data set was initially allocated, adjusting the target data set to accommodate the allocation change;

(m) updating a corresponding catalog and system meta-data describing the source data set to point to the target data set; and then

(n) signaling release to open the target data set.

Claim 47 was discussed in the interview. Many, if not all, of the steps distinguish over the prior art for reasons discussed above. For the foregoing reasons, all of the pending claims should now be allowed.

Respectfully submitted,

Mainstar Software Corporation

By _____

Micah D. Stolowitz
Registration No. 32,758
Direct Dial Telephone: 503-294-9189

STOEL RIVES LLP
900 SW Fifth Avenue, Suite 2600
Portland, OR 97204-1268
Telephone: (503) 224-3380
Facsimile: (503) 220-2480
Attorney Docket No.: 70205/11:1